

Filtering from Observations on Stiefel Manifolds ¹

J  r  mie Boulanger, Salem Said, Nicolas Le Bihan, and Jonathan H. Manton, *Senior member*

Abstract

This paper considers the problem of optimal filtering for partially observed signals taking values on the rotation group. More precisely, one or more components are considered not to be available in the measurement of the attitude of a 3D rigid body. In such cases, the observed signal takes its values on a Stiefel manifold. It is demonstrated how to filter the observed signal through the *anti-development* built from observations. A particle filter implementation is proposed to perform the estimation of the signal partially observed and corrupted by noise. The sampling issue is also addressed and interpolation methods are introduced. Illustration of the proposed technique on synthetic data demonstrates the ability of the approach to estimate the angular velocity of a partially observed 3D system partially observed.

I. INTRODUCTION

In numerous engineering problems, systems with states having values and evolving on the special orthogonal group $SO(n)$ can be encountered [1], [2], [3], [4], [5], [6]. In order to control such systems, their angular velocity must be estimated from possibly noisy measurements. This paper considers the case where only partial observations of the system are available, *i.e.* not all the components of the movement are recorded. The observation signal is modeled as a process taking its values on a Stiefel manifold. In addition, the presence of a multiplicative noise is considered in the observation process. Classical methods, including extended Kalman filter [7], [8] can not be applied directly here as they rely on the independent increments assumption. As explained later, it is not the case in the model we consider here. We propose to use the *anti-development signal* computed from the observed data. We present the way to build this signal, and address the sampling/interpolation issue as well. We also demonstrate how to perform optimal filtering on the anti-development signal. A numerical solution (particle filter) via a Monte-Carlo method is provided to perform this filtering and illustrated on the Stiefel manifold S^2 . The proposed technique is however valid for higher dimension Stiefel manifolds.

The rest of the paper is organized as follows. Section II presents the geometry of Stiefel manifolds based on the geometry of $SO(n)$ and the concept of horizontal space. Section III presents a time continuous theoretical solution to the filtering problem with observations in Stiefel manifolds. As opposed to the usual case, the noise cannot be considered additive anymore here in our model. The proposed solution is based on the antidevelopment, a defined with respect to the observation process that satisfies an additive noise model. Section ?? presents a theorem to overcome the problem of discrete sampling. Section V gives a practical solution based on a Monte-Carlo method for filtering. Finally Section VI considers the case of observation in $SO(n)$ and compare different approximation to the optimal solution.

II. GEOMETRY OF STIEFEL MANIFOLDS

The Stiefel manifold $V_{n,k}$ is the set of orthonormal k -frames in \mathbb{R}^n . It is well known and used in linear algebra to describe principal subspaces [9] and has found applications in sensors array [10], statistics [11], optimization

J. Boulanger is with the University of Rouen, France. email: jeremie.boulanger1@univ-rouen.fr.

S. Said is with the IMS, University of Bordeaux, France. email: salem.said@ims-bordeaux.fr

N. Le Bihan is with the CNRS, University of Melbourne, Australia. email: nicolas.le-bihan@gipsa-lab.grenoble-inp.fr. His research was supported by the ERA, European Union, through the International Outgoing Fellowship (IOF GeoSToSip 326176) program of the 7th PCRD.

J.H. Manton is with the University of Melbourne, Australia. email: jmanton@unimelb.edu.au

[12], channel estimation in wireless communications [13] or in light independent scene representation in computer vision [14].

First, recall that a $n \times n$ matrix R with real components is an element of the rotation group $SO(n)$ if it is orthogonal and has a unit determinant. This is to say that $R \in SO(n)$ iff:

$$R^T R = I_n \quad \text{and} \quad \det R = 1 \quad (1)$$

where I_n denotes the $n \times n$ identity matrix. Intuitively, $SO(n)$ is the set of positively oriented orthonormal basis vectors of \mathbb{R}^n .

In \mathbb{R}^n , the Stiefel manifold $V_{n,k}$ is defined as the set of matrices $P \in \mathbb{R}^{n \times k}$ such that:

$$P^T P = I_k \quad (2)$$

and with $k \leq n$. For example, if $k = 1$, then $V_{n,1}$ is the hypersphere S^{n-1} , i.e. the set of unit vectors in \mathbb{R}^n . If $k = n$, then $V_{n,n}$ corresponds to the orthogonal group $O(n)$.

Let $\Pi : SO(n) \rightarrow V_{n,k}$ be the projection consisting in the truncation of the $n - k$ last columns of a rotation matrix, and let us denote:

$$\Pi(R) = P \quad (3)$$

When $k \leq n - 2$, the projection is not injective. In this case, a matrix $P \in V_{n,k}$ can be completed by different sets of orthonormal vectors to form an oriented orthonormal basis of \mathbb{R}^n , which means that in such cases:

$$\Pi(R_1) = \Pi(R_2) \Leftrightarrow R_1 = R_2 \begin{pmatrix} I_k & 0 \\ 0 & C \end{pmatrix}$$

with $R_1, R_2 \in SO(n)$, $C \in SO(n - k)$.

However, if $k \leq n$, then Π is clearly surjective, i.e. $\Pi(SO(n)) \in V_{n,k}$ as the k first columns of a rotation matrix are orthonormal vectors. Therefore, $\Pi(R)^T \Pi(R) = I_k$ for $R \in SO(n)$. The Stiefel manifold $V_{n,k}$ can then be described as:

$$V_{n,k} = \{\Pi(R), R \in SO(n)\} \quad (4)$$

Note that the case $k = n$ needs special care. Indeed, $V_{n,n} = O(n)$ is the group of orthonormal matrix and is composed of two connected components: the set of orthonormal matrices with a positive determinant (positively oriented basis) $SO(n)$ and the set of orthonormal matrices with a negative determinant (negatively oriented basis). In this study, we will consider continuous random processes $P_t \in V_{n,k}$ which will solely belong to the same component of $V_{n,k}$ as their initial value P_0 belongs to. Therefore, if $\det(P_0) = +1$, then $\Pi = Id$ covers all the reachable points in the Stiefel manifold from $SO(n)$. If $\det(P_0) = -1$, considering $\Pi(R)$ as the application reversing the sign of the last column of R allows Π to cover all the reachable points in the Stiefel manifold from $SO(n)$. Consequently, expression (4) can be extended to the case where $k = n$, by considering only one connected component. This case will be considered in Section VI.

As $V_{n,k}$ can be constructed from $SO(n)$, we now investigate how the geometry of $V_{n,k}$ can be described using the geometry of $SO(n)$. From its definition, the projection Π is left invariant:

$$R_1 \Pi(R_2) = \Pi(R_1 R_2) \quad (5)$$

with $R_1, R_2 \in SO(n)$.

As Π is surjective, one also get the action of $SO(n)$ on $V_{n,k}$. If $R \in SO(n)$ and $P \in V_{n,k}$, then $RP \in V_{n,k}$. This property will be used later on to describe a process on $V_{n,k}$ via the action of $SO(n)$. This group action can be visualized by considering the example of the sphere $V_{3,1} \cong S^2$. Points at the surface of the sphere S^2 can reach all the locations on this manifold through the transitive action of $SO(3)$ on the sphere: $SO(3) \times S^2 \rightarrow S^2$.

First, let us identify the tangent bundle of $V_{n,k}$, denoted $TV_{n,k}$. It will be of use in Section III to define stochastic processes via the action of $SO(n)$ in the space tangent to a point in $V_{n,k}$. Denote $\mathfrak{so}(n)$ the Lie algebra¹ associated to the Lie group $SO(n)$ and let $\chi : \mathfrak{so}(n) \times V_{n,k} \rightarrow TV_{n,k}$ be the application defined by:

$$\chi(\sigma, P) = \left(\frac{d}{dt} \exp(t\sigma) P \right) \Big|_{t=0} = \sigma P \quad (6)$$

¹ $\mathfrak{so}(n)$ is the algebra of real-valued anti-symmetric matrices of size $n \times n$.

We can show by inclusion and dimension equality that $\chi(\cdot, P)$ is surjective onto $T_P V_{n,k}$, i.e. $T_P V_{n,k} = \{\sigma P, \sigma \in \mathfrak{so}(n)\}$, where we used the notation $T_P V_{n,k}$ for the tangent space attached to a point $P \in V_{n,k}$.

Now, for a given point $P \in V_{n,k}$, let $R \in SO(n)$ be a pre-image of P via Π , i.e. $P = \Pi(R)$. As Π is surjective, $\Pi^{-1}(P) \neq \emptyset$ and R is well defined. Then, the vertical space [14] at the point P , denoted \mathcal{V}_R , is defined as:

$$\mathcal{V}_R = \text{Ker } d\Pi_R \quad (7)$$

where $d\Pi_R$ is the differential of Π at the point R . By definition of $d\Pi_R$, the vertical space \mathcal{V}_R is a subspace of the tangent space $T_R SO(n)$. In the case when $n = 3$ and $k = 1$, then $V_{3,1} \cong S^2$ and P is a point on the unit sphere in \mathbb{R}^3 . The vertical space corresponds to the set of rotations which have their axis aligned with P . Such rotations leave P invariant. Figure 1 displays a graphical interpretation of the vertical space \mathcal{V}_R .




Fig. 1. Graphical representation of the vertical space \mathcal{V}_R at P for the case $P \in V_{3,1}$. For a rotation R acting on P , \mathcal{V}_R is the orthogonal complement of \mathcal{H}_R in $T_P V_{3,1}$. If the axis of the rotation of R is parallel to P , then its action is in \mathcal{V}_R and it is not visible as P is rotating about itself.




Fig. 2. Graphical representation of the horizontal space \mathcal{H}_R at P for the case $P \in V_{3,1}$. For a rotation R acting on P , \mathcal{H}_R is a subspace of $T_P V_{3,1}$. If the axis of rotation of R is orthogonal to P , then its action is in \mathcal{H}_R and it is visible.

Making use of the standard scalar product on $\mathfrak{so}(n)$ which reads for any skew-symmetric matrices $\sigma, \varsigma \in \mathfrak{so}(n)$ like:

$$\langle \sigma, \varsigma \rangle = \frac{1}{2} \text{tr}(\sigma^T \varsigma) \quad (8)$$

it is possible to construct the orthogonal complement of \mathcal{V}_R in $T_R SO(n)$, called the *horizontal space* and denoted \mathcal{H}_R . We have then that:

$$T_R SO(n) = \mathcal{V}_R \oplus \mathcal{H}_R \quad (9)$$




Fig. 3. Illustration of the different notions introduced to describe the Stiefel manifold $V_{n,k}$ as the image from the projection Π of $SO(n)$. The horizontal \mathcal{H}_R and vertical \mathcal{V}_R spaces are dependent of the chosen pre-image but the translation into $\mathfrak{so}(n)$ via ω is invariant with respect to the choice of the pre-image.

Figure 2 displays a graphical interpretation of the horizontal space \mathcal{H}_R .

As $d\Pi_R$ is linear, the restriction of $d\Pi_R$ to \mathcal{H}_R is bijective. In other words, $T_P V_{n,k}$ and \mathcal{H}_R are isomorphic. For a vector $v \in T_P V_{n,k}$, let $v^{\mathcal{H}} \in T_R SO(n)$ be the vector defined as:

$$d\Pi_R(v^{\mathcal{H}}) = v \quad (10)$$

For example, consider again the case of the Stiefel manifold $V_{3,1} \cong \mathcal{S}^2$. Considering $P = (1, 0, 0)^T \in V_{3,1}$, the matrices $R_1 = (e_1, e_2, e_3)$ and $R_2 = (e_1, e_3, -e_2)$ for $\{e_i\}_{i \leq 3}$ the canonical basis of \mathbb{R}^3 are both pre-images of P , i.e. $\Pi(R_1) = \Pi(R_2) = P$. The application $\chi(\cdot, P)$ describes the tangent space $T_P V_{3,1}$ like $\chi(\sigma, P) = (0, \sigma_{21}, \sigma_{31})^T$ where σ_{ij} is the (i, j) matrix elements of $\sigma \in \mathfrak{so}(3)$. At the pre-image R_1 , the vertical and horizontal spaces are thus defined as:

$$\mathcal{V}_{R_1} = \left\{ \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & \alpha \\ 0 & -\alpha & 0 \end{pmatrix}, \alpha \in \mathbb{R} \right\}$$

and:

$$\mathcal{H}_{R_1} = \left\{ \begin{pmatrix} 0 & \beta & \gamma \\ -\beta & 0 & 0 \\ -\gamma & 0 & 0 \end{pmatrix}, \beta, \gamma \in \mathbb{R} \right\}.$$

At the pre-image R_2 , the vertical and horizontal spaces are defined as:

$$\mathcal{V}_{R_2} = \left\{ \begin{pmatrix} 0 & 0 & 0 \\ 0 & \alpha & 0 \\ 0 & 0 & \alpha \end{pmatrix}, \alpha \in \mathbb{R} \right\}$$

and:

$$\mathcal{H}_{R_2} = \left\{ \begin{pmatrix} 0 & \beta & \gamma \\ \gamma & 0 & 0 \\ -\beta & 0 & 0 \end{pmatrix}, \beta, \gamma \in \mathbb{R} \right\}.$$

One can directly check that spaces $\mathcal{H}_{R_1} R_1^T$ and $\mathcal{H}_{R_2} R_2^T$ are identical. This is true even if the horizontal subspaces \mathcal{H}_{R_1} and \mathcal{H}_{R_2} are different, and is a consequence of the fact that they are defined by a different pre-image of P . A graphical illustration of the notion of horizontal and vertical spaces is displayed in figure 3

Due to the isomorphism between $T_P V_{n,k}$ and \mathcal{H}_R , the application $\chi(\cdot, P)$ restricted to $\mathcal{H}_R R^T \in \mathfrak{so}(n)$ with $\Pi(R) = P$ is bijective. In other words, $\chi(\cdot, P)|_{\mathcal{H}_R}$ is invertible. Let $\omega : T V_{n,k} \rightarrow \mathfrak{so}(n)$ denote this inverse and let us call it the *restricted inverse*. It then reads:

$$\omega(v, P) = v^{\mathcal{H}} R^T. \quad (11)$$

The term $v^{\mathcal{H}}$ is the horizontal vector from the tangent space to R . Despite the definition of ω being dependent on R , this is not the case because $v^{\mathcal{H}}$ also depends on R , and, in the end, the term $v^{\mathcal{H}} R^T$ is independent of R . Finally,

it is possible to define a metric on $V_{n,k}$ using the metric on $SO(n)$. Let $\langle \cdot, \cdot \rangle_P$ be the metric defined as:

$$\langle v_1, v_2 \rangle_P = \langle v_1^{\mathcal{H}}, v_2^{\mathcal{H}} \rangle_R \quad (12)$$

for any two vectors $v_1, v_2 \in T_P V_{n,k}$, with $\Pi(R) = P$, and where $\langle \cdot, \cdot \rangle_R$ denotes the scalar product defined in $T_R SO(n)$.

III. FILTERING FROM OBSERVATIONS ON STIEFEL MANIFOLDS

We consider the problem of a partially observed system whose state evolves on the rotation group $SO(n)$. In practice, such observations may come from flawed sensors or devices, leading to the availability of a limited part of the signal to filter. For example, in the context of satellite's control, existing algorithms require the knowledge of the angular velocity and the orientation of the satellite to monitor its orientation [2][1]. This angular velocity is determined from different internal sensors. However, if some of these sensors become faulty, the velocity of the satellite is no more available and the satellite cannot be controlled properly anymore.

The presented algorithm proposes to tackle the problem of lack in parts of the signal to filter and takes advantage of the available observations to perform optimal filtering. More precisely, we present a technique to obtain an estimate of the velocity of the system with only partial observations of its orientation, *i.e.* partial observations on $SO(n)$.

A. Observation model

The model considered is as follows: a process $S_t \in SO(n)$ is defined by its angular velocity $x_t \in \mathfrak{so}(n)$ where t represents time. Our aim is to obtain an estimate of the angular velocity x_t based on observations of S_t which are not complete as well as noisy. The process x_t is here assumed to be the solution of the following linear stochastic differential equation in $\mathfrak{so}(n)$:

$$dx_t = Fx_t + db_t \quad (13)$$

where $b_t \in \mathfrak{so}(n)$ is a Brownian motion with variance σ_b^2 . In this case, x_t is a Markov process and its transition kernel for time $t + s$ based on x_s is denoted $q_t(x_s, \cdot)$.

The partial observation is here modeled as a process P_t on the Stiefel manifold $V_{n,k}$, *i.e.* only k components of S_t amongst the total of n components are known. The filtering problem then reads: we want to estimate $x_t \in \mathfrak{so}(n)$ from $P_t \in V_{n,k}$ defined as $P_t = \Pi(S_t)$ in the presence of noise. The noise is modeled by a Brownian motion $w_t \in \mathfrak{so}(n)$ with variance σ_w^2 independent from x_t acting in $\mathfrak{so}(n)$. As x_t is the angular velocity of the observed processed P_t , then P_t is solution of the stochastic differential equation:

$$dP_t = (x_t dt + \circ dw_t) P_t \quad (14)$$

where notation \circ is used to denote the Stratonovich integral.

Due to the presence of the noise w_t , x_t cannot be exactly determined. Instead, we want to determine the distribution π_t of x_t conditioned by the observation of $\mathcal{P}_s = \{P_s, s \leq t\}$. It is possible to construct some estimator for x_t based on its conditional distribution π_t .

It is noticeable that despite that the noise acts additively in the tangent space $T_{P_t} V_{n,k}$, it acts as a multiplicative noise for the process P_t , preventing us from using usual filtering methods. Indeed, classical methods like Kalman filter rely on the independence of the increments dP_t . However, this is not applicable in our case as the increments depends of P_t .

B. The anti-development solution

We propose a solution based on the concept of anti-development. It consists in constructing a process z_t in one-to-one correspondence with P_t such that z_t is solution of a stochastic differential equation with additive noise. The likelihood used to compute the solution is then based on z_t .

Let $z_t \in \mathfrak{so}(n)$ and $R_t \in SO(n)$ be defined as:

$$\begin{aligned} dz_t &= \omega(\circ dP_t, P_t) \\ dR_t &= (\circ dz_t) R_t \end{aligned} \quad (15)$$

where ω is as defined previously in (11), the *restricted inverse* of χ .

The process z_t is called the anti-development of P_t and R_t is called the horizontal lift of P_t [15]. An illustrative example of the anti-development on $V_{3,1} = S^2$ is presented in figure 4.

The process z_t is the accumulation of the increments in the tangent space $T_P V_{n,k}$ whereas R_t is the rotational process constructed by considering that the component in the vertical space is null. These processes are equivalent, in terms of information to P_t because P_t can be constructed like $P_t = R_t R_0^T P_0$ or $dP_t = \chi(dz_t, P_t)$. Conditioning the distribution of x_t by the observation of P_t is equivalent as conditioning by the observation of z_t . However, the anti-development is a solution, as opposed to P_t , of a stochastic differential equation with additive noise. It should be noticed that in general cases, $S_t \neq R_t$. Despite the vertical component has no action in the Stiefel manifold $V_{n,k}$, it still has some effect in $SO(n)$. This involves that in general, $S_t \neq R_t$.

Replacing dP_t in (15) by its expression from (14) gives:

$$\begin{aligned} dz_t &= \omega(dP_t, P_t) \\ &= \omega((x_t dt + \odot dw_t), P_t, P_t) \\ &= \omega(\chi(x_t dt + \odot dw_t), P_t, P_t) \\ &= \omega(H_t, P_t) + \odot d\beta_t \end{aligned} \quad (16)$$

where $H_t = \chi(x_t, P_t)$ and $d\beta_t = \omega(\chi(\odot dw_t, P_t), P_t)$. By definition of $d\beta_t$, the process β_t is constructed from the k first components of w_t . Therefore, β_t is a Brownian process with a variance that can be diagonalized as $\sigma_w^2 I_{n,k}$ where $I_{n,k} = \text{diag}(1, \dots, 1, 0, \dots, 0)$ with k non-zero elements.

This way, our filtering problem from observations in the Stiefel manifold $V_{n,k}$ with multiplicative noise is now reduced to a filtering problem in $\mathfrak{so}(n)$ with additive noise. For a test function ϕ , we want to determine $\pi(\phi) = \mathbb{E}[\phi(x)|\mathcal{P}_t]$. The solution is therefore given by applying usual filtering methods [16] to the anti-development z_t defined in (15):

$$\pi_t(\phi) = \frac{\rho_t(\phi)}{\rho_t(1)}, \quad (17)$$

with $\rho_t(\phi) = \mathbb{E}[\phi(x') L_t(P, x') | \mathcal{P}_t]$, where x'_t is a copy of x_t independent of P_t and $\mathcal{P}_t = \{P_s, s \leq t\}$. The likelihood L_t is defined as:

$$L_t(P, x') = \exp \left(\frac{1}{\sigma_w^2} \int_0^t \langle x'_s, dz_s \rangle - \frac{1}{2} \|x'_s\|^2 ds \right) \quad (18)$$

with $dz_t = \omega(\odot dP_t, P_t)$. By definition of the inner product in (12), the likelihood in (18) can be rewritten like:

$$L_t(P, x') = \exp \left(\frac{1}{\sigma_w^2} \int_0^t \langle H'_s, dP_s \rangle - \frac{1}{2} \|H'_s\|^2 ds \right) \quad (19)$$

with $H'_s = \chi(x'_s, P_s)$ a copy of H_s in the distribution sense. Expression (19) is more amenable than the one from equation (18) as it does not require the computation of R_t . The integrand can directly be determined from the observations without constructing any auxiliary process. However, using expression (19), the model for H_t is not linear, even if x_t is the solution of a linear model. As a consequence, expressing $\rho_t(\phi)$ is a complicated task. Nevertheless, it is still possible to get an approximation of the solution, using numerical methods, like the particle filter for example. Before proposing a filtering solution, we address the issue due to the discrete nature of the observation of P_t .

IV. INTERPOLATION FUNCTION

It must be noted that the likelihood function L_t given in (19) requires the full observation of the process $\{P_s\}_{s \leq t}$ to compute the integrand. In practice, it is not possible to have a continuous observation of P_t . Only discrete samples are available. Let δt be the sampling period. Between two samples, P_t must be approximated using an interpolation function. This interpolation function must be chosen to minimize the approximation error as a function of the sampling period.




Fig. 4. Example of a trajectory of P_t (red on the sphere) on $V_{3,1} = S^2$ at three successive times (Top to Bottom). The anti-development z_t is displayed in red in the plane over the sphere. It can be obtained by considering the trace left by the sphere when rolling without slipping on the plane, and with P_t as a contact point. The anti-development is solution of a stochastic differential equation with additive noise, as opposed to P_t .

Let $\text{Int} : V_{n,k} \times V_{n,k} \rightarrow \mathfrak{so}(n)$ be an interpolation function. It is thus required that Int should be such that given $\delta z_k = \text{Int}(P_{k\delta t}, P_{(k+1)\delta t})$, the likelihood based on discrete observation will converge to the continuous solution for $\delta t \rightarrow 0$.

Theorem 1: The Riemann sum

$$\tilde{\mathcal{S}}_{n\delta t} = \sum_{k=0}^n \langle x_{k\delta t}, \text{Int}(P_{k\delta t}, P_{(k+1)\delta t}) \rangle$$

with $n = t/\delta t$ converges towards

$$\mathcal{S}_t = \int_0^t \langle x_s, dz_s \rangle$$

in the sense $\mathbb{E}[|\tilde{\mathcal{S}}_{n\delta t} - \mathcal{S}_t|^2] \rightarrow 0$ when $\delta \rightarrow 0$ if the interpolation function $\text{Int} : V_{n,k} \times V_{n,k} \rightarrow \mathfrak{so}(n)$ satisfies the following conditions:

- Its diagonal elements are nul, i.e. $\text{Int}(P, P) = 0$ for all $P \in V_{n,k}$.
- The function Int is $\mathcal{C}^2(V_{n,k})$.
- $\nabla \text{Int}(P, P)[v] = v^{\mathcal{H}} R^T$ for all $v \in T_P V_{n,k}$ and $\Pi(R) = P$.
- $\nabla^2 \text{Int}(P, P)[v] = 0$ for all $v \in T_P V_{n,k}$.

Where the differentials (∇I and $\nabla^2 I$) are computed with respect to the second variable.

Proof 1: Considering the function $f = \text{Int}(P_{k\delta t}, \cdot)$, then condition 2) allows the use of Itô lemma:

$$\begin{aligned} f(P_{(k+1)\delta t}) - f(P_{k\delta t}) &= \int_{k\delta t}^{(k+1)\delta t} \nabla f[dP_s] + \\ &\quad \frac{1}{2} \int_{k\delta t}^{(k+1)\delta t} \text{trace}([dP_s]^T \nabla^2 f[dP_s]) \end{aligned}$$

Condition 4) sets the last term to $o(\delta t)$ whereas condition 3) sets the first term to $\int_{k\delta t}^{(k+1)\delta t} (dP_s)^{\mathcal{H}} R_s^T + o(\delta t)$. Then, replacing f by $\text{Int}(P_{k\delta t}, \cdot)$

$$\begin{aligned} &\text{Int}(P_{k\delta t}, P_{(k+1)\delta t}) - \text{Int}(P_{k\delta t}, P_{k\delta t}) \\ &= \int_{k\delta t}^{(k+1)\delta t} (dP_s)^{\mathcal{H}} R_s^T \end{aligned}$$

Now, thanks to condition 1), one gets that $\text{Int}(P_{k\delta t}, P_{k\delta t}) = 0$.

Separating the integral \mathcal{S}_t into $t/\delta t$ short integrals gives, up to a remaining integral between t and $kt/\delta t$, the following expression:

$$\begin{aligned} &\mathbb{E} \left[|\tilde{\mathcal{S}}_{n\delta t} - \mathcal{S}_t|^2 \right] \\ &= \mathbb{E} \left[\left| \sum_{k \leq t/\delta t} \int_{k\delta t}^{(k+1)\delta t} \langle x_{k\delta t} - x_s, dz_s \rangle \right|^2 \right] \\ &= \mathbb{E} \left[\left| \sum_{k \leq t/\delta t} \int_{k\delta t}^{(k+1)\delta t} \langle x_{k\delta t} - x_s, (dR_s) R_s^T \rangle \right|^2 \right] \\ &\quad (\text{as } \sigma_w^2 I_n \text{ is orthogonal to } \mathfrak{so}(n)) \\ &\leq \sum_{k \leq t/\delta t} \mathbb{E} \left[\left| \int_{k\delta t}^{(k+1)\delta t} \langle \hat{x}_{k\delta t} - \hat{x}_s, (dP_s)^{\mathcal{H}} R_s^T \rangle \right|^2 \right] \end{aligned}$$

The square variations of x_t during a time δt being bounded by $O(\delta t)$, the variation of x_t are bounded by $O(\delta t)$ too. The integral is then bounded by $O(\delta t^2)$ thanks to the Itô isometry property of the integral

$$\begin{aligned} & \mathbb{E} \left[|\tilde{\mathcal{S}}_{n\delta t} - \mathcal{S}_t|^2 \right] \\ &= \mathbb{E} \left[\sum_{k \leq t/\delta t} \int_{k\delta t}^{(k+1)\delta t} O(\delta t) \|(dP_s)^{\mathcal{H}} R_s^T\|^2 \right] \\ &= \sum_{k \leq t/\delta t} O(\delta t^2) \\ &= O(\delta t) \end{aligned}$$

This means that S^δ converges towards \mathcal{I} in the mean square error sense when δt shrinks to 0 with a linear convergence rate.

Based on this theorem, we can show that a linear interpolation between successive samples can be used to approximate the likelihood.

Proposition 1: The interpolation function $\text{Int}(P', P) = \frac{1}{2} d\Pi_{R'} (RR'^T - R'R^T)$ with $\Pi(R) = P$ and $\Pi(R') = P'$ satisfies the conditions of Theorem 1.

Proof 2: The conditions 1) and 2) from theorem 1 are directly verified. Using the notation $\sigma = v^{\mathcal{H}} R^T$, Condition 3) is verified as:

$$\begin{aligned} \nabla \text{Int}(P, P)[v] &= \frac{d}{dt} \text{Int}(P, \exp(t\sigma)P)|_{t=0} \\ &= \frac{d}{dt} \frac{1}{2} (\exp(t\sigma)RR^T - R^T \exp(t\sigma)^T) |_{t=0} \\ &= \frac{d}{dt} \frac{1}{2} (\exp(t\sigma) + \exp(t\sigma)) |_{t=0} \\ &\text{as } \sigma \in \mathfrak{so}(n) \\ &= \sigma = vR^T \end{aligned}$$

Finally, condition 4) is also satisfied because:

$$\begin{aligned} & \nabla^2 \text{Int}(P, P)[v] \\ &= \frac{d^2}{dt ds} \text{Int}(P, \exp(s\sigma) \exp(t\sigma)P) |_{t=0, s=0} \\ &= \frac{d}{dt} \frac{1}{2} (\exp(s\sigma) \exp(t\sigma)RR^T - \\ &\quad R^T \exp(t\sigma)^T \exp(s\sigma)^T) |_{t=0} \\ &= \frac{d}{dt} \frac{1}{2} (\exp((t+s)\sigma) - \exp((t+s)\sigma)) |_{t=0} \\ &= 0 \end{aligned}$$

Thanks to Proposition 1, we are now able to implement a Monte-Carlo solution based on discrete samples. For a better readability, the interpolation term between two successive samples $\text{Int}(P_{k\delta t}, P_{(k+1)\delta t})$ will be denoted δP_k in the sequel.

V. PRACTICAL SOLUTION TO THE FILTERING PROBLEM

A. Implementation via a Monte-Carlo method

The particle filter is a Monte-Carlo method to approximate the solution given by Equation (17) for non linear model. Despite that particle filters has been heavily used and studied [17], the application of this method to perform estimation from partial observation on the Stiefel manifold has not been used before.

The main idea of the particle filter is to approximate the expectation in (17) using the law of large numbers. Recall that:

$$\rho_t(\phi) = \mathbb{E}[\phi(x')L_t(P, x')|\mathcal{P}_t]$$

with L_t the likelihood defined in Equation (19). The process x'_t is a copy of x_t (in the sense with the same model of propagation) but, contrary to x_t , independent from P_t .

Let X_t^i , with $i = 1, \dots, N$, be N processes identical to x_t called particles. They represent candidates to estimate the process x_t . The law of large numbers states that ρ_t^N defined as:

$$\rho_t^N(\phi) = \frac{1}{N} \sum_{i=1}^N \phi(X_t^i) L_t(P, X_t^i)$$

will converge with N almost surely to $\rho_t(\phi)$. Therefore, the previous equation gives an approximation of the solution by determining ρ_t^N and normalizing it. It is noticeable that the particles are not observable and must be simulated. This means that the model of propagation (13) for x_t should be known.

Furthermore, it is assumed that the process P_t is not continuously observed and let denotes δt the sampling time. These means that it is necessary to consider a time discretized version of ρ_t^N , denoted ρ_n^N , with $n = \lfloor t/\delta t \rfloor$, as:

$$\rho_n^N(\phi) = \sum_i \phi(\tilde{X}^i) L_{\delta t}(\tilde{P}_n, \tilde{X}_0^i, \dots, \tilde{X}_n^i)$$

where $\tilde{X}_n^i = X_{n\delta t}^i$, $\tilde{P}_n = \{P_{k\delta t}, k \leq n\}$ and the likelihood $L_{\delta t}$ is defined by:

$$L_{\delta t}(i, n) = \exp \left(\frac{1}{\sigma_w^2} \sum_{k \leq n} \langle H_k^i, \delta \tilde{P}_k \rangle - \frac{1}{2} \|H_k^i\|^2 \right)$$

with $H_k^i = \chi(X^i, \tilde{P}_k)$. It has been proved in Proposition 1 that the discrete likelihood $L_{\delta t}$ converges towards the likelihood L_t from (19).

In order to implement such a solution, two independent problems must be tackled:

- The simulation of the particles \tilde{X}^i :

It will be supposed that x_t is a Markov process. Consequently, \tilde{X}^i is a Markov chain with transition kernel $q_{\delta t}$ defined after (13) and \tilde{X}_{n+1}^i is directly sampled from $q_{\delta t}(\tilde{X}_n^i, \cdot)$.

- The computation of the likelihood $L_{\delta t}(\tilde{P}_n, \tilde{X}_0^i, \dots, \tilde{X}_n^i)$:

Considering the last term of the Riemannian sums:

$$\begin{aligned} L_{\delta t}(i, n) &= \exp \left(\frac{1}{\sigma_w^2} \sum_{k \leq n} \langle H_k^i, \delta \tilde{P}_k \rangle - \frac{1}{2} \|H_k^i\|^2 \right) \\ &= L_{\delta t}(i, n-1) l_{\delta t}^i \end{aligned}$$

where $l_{\delta t}^i = \exp \left(\frac{1}{\sigma_w^2} \langle H_n^i, \delta \tilde{P}_n \rangle - \frac{1}{2} \|H_n^i\|^2 \right)$, this decomposition shows that the likelihood can be computed adaptively when new samples are available.

This leads to Algorithm 1, here described to estimate the conditional distribution $\pi_n^{N, \delta t}(\phi)$ as

$$\pi_n^{N, \delta t}(\phi) = \frac{\rho_n^{N, \delta t}(\phi)}{\rho_n^{N, \delta t}(1)} = \sum_{i=1}^N \phi(\tilde{X}_n^i) w_n^i$$

where the coefficient

$$w_n^i = \frac{L_{\delta t}(\tilde{P}_n, \tilde{X}_0^i, \dots, \tilde{X}_n^i)}{\sum_{j=1}^N L_{\delta t}(\tilde{P}_n, \tilde{X}_0^j, \dots, \tilde{X}_n^j)}$$

is called the "weight" associated to the particle i .

The normalization step (step 3) is not only here to compute $\pi_n^{N, \delta t}$ instead of $\rho_n^{N, \delta t}$ but also to numerically stabilize the computation of the weights. As they are usually smaller than 1, their consecutive multiplications lead to small values.

Algorithm 1 Particle filter algorithm

- For the initialization, generate N particles from a priori p_0 : $\tilde{X}_0^i \sim p_0$ and set $w_0^i = 1/N$.
- At a time $n > 0$:

- 1) Propagate the particles $\tilde{X}_n^i \sim q_{\delta t}(\tilde{X}_{n-1}^i, \cdot)$
- 2) Update the weight w_n^i of each particle as: $w_n^i = w_{n-1}^i l_{\delta t}^i$ with

$$l_{\delta t}^i = \exp\left(\frac{1}{\sigma_w^2} \langle H_n^i, \delta \tilde{P}_n \rangle - \frac{1}{2} \|H_n^i\|^2\right).$$

- 3) Normalize the weights: $w_n^i = w_n^i / \sum_j w_n^j$
- 4) If $(\sum_i (w_n^i)^2)^{-1} < N/2$, generates $[m^1 \dots m^N] \sim \text{multinomial}(w^1 \dots w^N)$ such that $\sum_i m^i = N$. Then, clones \tilde{X}_n^i m^i -times and set $w_n^i = \frac{1}{N}$.
- 5) Estimate $\pi_t(\phi)$ with

$$\pi_n^{N, \delta t}(\phi) = \sum_i \phi(\tilde{X}_n^i) w_n^i.$$

Step 4 is called resampling. It is here to prevent a degeneracy due to the finite number of particles. Indeed, the particles are propagating without any restriction or drift imposed by the observation. Without resampling, particles would just explore the space and as they tend to drift away from x_t , they would become a bad approximation of x_t , because the mean square error $\mathbb{E}[(X_t^i - x_t)^2]$ is linearly growing with time. The number of particles being fixed, their weights quickly degenerate as they are diffusing away from x_t . Due to the normalization step, this leads to the concentration of all the ponderation into one single particle. Even if this particle is the best candidate amongst the all the particles, the mean square error is still linearly growing.

The resampling step consists in killing the particles far away from x_t (in fact, killing the particles with low weights) and cloning the remaining ones. In order to measure if the particles are scattered away from x_t , one commonly used criteria is a threshold based on the Effective Sample Size (ESS_w) defined as

$$ESS_w = \left(\sum_i (w_n^i)^2 \right)^{-1}.$$

When ESS_w is lower that say $N/2$, then particles need to be resampled (the criteria ESS_w is small when only a few particles have a preponderant weight). To resample the particles, one can for example sample $[m^1 \dots m^N]$ from a multinomial distribution

$$[m^1 \dots m^N] \sim \text{multinomial}(w^1 \dots w^N)$$

such that $\sum_i m^i = N$ (to keep the number of particles constant) and clone the i^{th} particle m^i times. If w^i is high (particle with a good likelihood, thus a good candidate), then m^i should be high too. This effect will tend to keep only the good candidates, based on the likelihood. However, instead of resampling when the Effective Sample Size becomes too low, resampling is made after a fixed given time. In fact, resampling can be realized at every iteration but it is, computationally speaking, expensive and does not bring noticeable improvements [17].

B. Simulation results

This subsection describes the results obtained from a numerical implementation of the particle filter detailed in Algorithm 1. For this simulation, the chosen Stiefel manifold was the sphere $V_{3,1} = S^2$. The process $x_t \in \mathbb{R}^3$ is a Brownian motion with a unit variance. The variance of the noise is fixed to $\sigma_w^2 = 1$.

To approximate x_t , $N = 500$ particles are generated from a normal prior distribution p_0 centered around the origin with a variance of 2. Using more particles does not significantly improve the results. Considering a bad prior for generating the particles is not a big issue as the resampling step quickly eliminates the wrong candidates for the estimation. In the first two cases displayed in Figure 5 (Top and Bottom), the time step for the observation is $\delta t = 10^{-2}s$ and it is set to $\delta_t = 0.5s$ for the last case (Bottom). The time step for creating the simulation has been fixed to $10^{-3}s$, which is sufficient to consider the process continuous with respect to the observation time step.




Fig. 5. Evolution of the estimation (red) of each component of x_t (black), namely x_1 , x_2 and x_3 . (Top) The model used for x_t is a stair function and the times at which changes occur are known. The sampling time is $\delta t = 10^{-2}$. The particle algorithm properly converges. (Middle) x_t is a Brownian motion in \mathbb{R}^3 with unit variance. The algorithm can still estimate properly as the variations are not too fast. (Bottom) For the same model, the sampling time is reduced to $\delta t = 0.5$ and the filter is not able to track x_t anymore. The component in the vertical space changes too quickly.




Fig. 6. Projection of the particles (yellow), the state to estimate (blue) and the estimation (green) on the plane Ox, Oy and Ox, Oz . At the time of the snapshot, $P_t \approx [0, 1, 0]^T$. Consequently, the component along Oy cannot be estimated. The particles are distributed within an ellipse whose large axis is directed along Oy .

Figure 5 illustrates the results obtained for the estimation when the state x_t is a stair function (Top), then when x_t varies slowly (Middle) and finally when the sampling period is too large to be able to track properly the evolution of x_t (Bottom). As long as x_t is slowly varying with respect to δt , the algorithm is able to completely estimate x_t . In the case where x_t follows a stair function model, one could use a classical algorithm to detect abrupt changes in x_t in order to estimate the time instants where the particles should be sampled [18] (this was not effectively implemented in the results presented in the Top figure of 5, where it was simply assumed that the time where changes occur were known). When a change is detected, the particles are once again sampled from the initial priori to converge toward the new value. In the presented case (Top of figure 5), the particles will not drift away because they are at a constant position (they propagate with the same model as x_t). Particles strongly merge when they are resampled, leaving less and less possible choice.

It was mentioned earlier that the vertical component could not be estimated. However, the vertical space is defined with respect to the observation point P_t . As P_t will evolve on $V_{n,k}$, the vertical space will change too and the component on the initial vertical space can thus be estimated. As a consequence, it is finally possible to completely estimate the angular velocity. Note that in the case where x_t evolves slowly, the vertical component can still be estimated.

Now, if the angular velocity is evolving too quickly (with respect to the amount of time particles need to converge), it will not be possible to estimate the vertical component of x_t from the observation of P_t . The particles are distributed within an ellipse (see Figure 6, Top), whose large axis coincides with the direction of P_t . Along the direction of P_t , the estimation of x_t , the empirical average of the particles (yellow dots), is not satisfactory. However, in the other direction, the estimation is correct. This is due to the fact that the particles can only track

the component of x_t that has an impact on P_t . As the innovation term in Equation (19) is $\langle H_s, dP_s \rangle$, only the horizontal component can be observed and therefore estimated. Consequently, particles propagating along the vertical subspace (wich is, for S^2 , the line defined by P_t) are not penalized (their weight does not decrease) and are still considered as good candidates. As a consequence, the estimate is correct in the horizontal direction, but not in the vertical one. This last comment can be understood as highlighting the cases where the proposed particle filter failed at estimating correctly the complete set of components of x_t due to the lack of information in the observation.

In the next section, we consider a special case of our problem, namely when observations are complete.

VI. SPECIAL CASE FOR OBSERVATION FROM $SO(n)$

A. Optimal filtering in $SO(n)$

The special case when observations are in $V_{n,n}$ can be described using the technique presented in previous sections. However, the fact that P_t is in one of the continuous component of $O(n)$ allows us to treat it also in a different manner. As it is described in Section II, the process P_t can be considered like a process with values in $SO(n)$ (which is one of the two continuous components of $O(n)$) without any loss of generality.

As $P_t \in SO(n)$, then the map Π is the identity map and the application χ is invertible. Therefore, there is no vertical space over P_t and the tangent space is simply the horizontal space. Equation (16) then reads:

$$dz_t = x_t dt + \odot dw_t \quad (20)$$

where $dz_t = (\odot dP_t)P_t^T$. In this case, a numerical method is no more required as the full process is observed, *i.e* $H_t = x_t$. As the increments in Equation (20) are independent, and the noise is additive, a classic Kalman filter can be used where the anti-development z_t replaces the observed process P_t . The conditional distribution of x_t is then a Gaussian distribution with mean μ_t and a variance V_t such that:

$$\begin{aligned} d\mu_t &= F\mu_t dt + \frac{1}{\sigma_w^2} V_t (dz_t - \mu_t dt) \\ dV_t &= FV_t + V_t F^T - \frac{1}{\sigma_w^2} V_t^2 + \sigma_b^2. \end{aligned} \quad (21)$$

This solution represents the optimal filter for observation in $SO(n)$. However, in practice, the same issue as in Section III occurs due to the discretization of the observation.

B. Implementation of the solution

The discrete nature of the observation of P_t does not allow to continuously determine z_t . An approximation must be performed using an interpolation function. The linear interpolation function $\text{Int}(P', P) = \frac{1}{2} d\Pi_{R'}(PP'^T - P'P^T)$ still converges towards the continuous solution as the sampling period δt shrinks to 0. Amongst all the possible functions, one can also choose to use $\text{Int}(P', P) = \log(PP'^T)$ as it satisfies the conditions of Theorem 1.

Proposition 2: The interpolation function $\text{Int}(P', P) = \log(PP'^T)$ satisfies the conditions of Theorem 1.

Proof 3: The conditions 1) and 2) are direct. For condition 3), an element $v \in T_P SO(n)$ is described with right invariant vector fields $V = \sigma P$ with $\sigma \in \mathfrak{so}(n)$.

$$\begin{aligned} \nabla \text{Int}(P, P)[v] &= \frac{d}{dt} \text{Int}(P, \exp(t\sigma)R)|_{t=0} \\ &= \frac{d}{dt} \log(\exp(t\sigma)P^T)|_{t=0} \\ &= \frac{d}{dt} t\sigma|_{t=0} \\ &= \sigma = vP^T \end{aligned}$$




Fig. 7. Comparison of estimation results (MSE) between linear interpolation, geodesic interpolation, and classic Kalman filter with $\delta t = 0.01s$. The state to estimate is constant and the variance of the observation is $\sigma_w^2 = 1$

Condition 4) is also satisfied as:

$$\begin{aligned}
 & \nabla^2 \text{Int}(P, P)[v] \\
 &= \frac{d^2}{dt ds} \text{Int}(P, \exp(s\sigma) \exp(t\sigma) P) |_{t=0, s=0} \\
 &= \frac{d^2}{dt ds} \log(\exp((t+s)\sigma) P^T) |_{t=0} \\
 &= \frac{d^2}{dt ds} (t+s)\sigma |_{t=0, s=0} \\
 &= 0
 \end{aligned}$$

The function $\text{Int}(P', P) = \log(P P'^T)$ is called geodesic interpolation. This interpolation function has a higher computational cost than the linear interpolation one, but it is also more accurate for approximating the incremental term dz_t .

In order to illustrate this point, Figures 7 and 8 present the difference between a Kalman filter directly applied on the discrete observation of z_t (which is an optimal filter), and the different methods of interpolation based on the discrete observation of P_t .

Because the equation of the variance V_t does not depend on the observation and is just an isolated differential equation (with respect to the innovation), the variance for each algorithm used in Figure 7 is the same (considering that the initial value V_0 has always been chosen with the same value) and is consequently not displayed for this kind of comparison.

In the presented simulations, the process x_t is considered constant. The use of another model might not imply any significant change. The process z_t is first generated with a small time step ($10^{-4}s$) and with a variance $\sigma_w^2 = 1$ via an Euler scheme from the stochastic differential equation (20).

At the same time, a rotational process P_t is constructed from z_t . The construction of z_t and P_t is then realized with a time step small enough to consider them as "time continuous".

The process P_t is then sampled with a time step δt with $\delta t \gg 10^{-4}$. The performances of the filter from (21) with different methods of interpolation (linear and the geodesic) to approximate $\delta z_{k\delta t}$ are displayed in figure 7. In parallel, as z_t has been continuously generated, these performances are compared to a classical Kalman filter taking directly $\delta z_{k\delta t} = z_{(k+1)\delta t} - z_{k\delta t}$. Recall that this term is not available due to the discrete observation of P_t . Knowing that in the case of additive noise, the Kalman filter is optimal [16] and because the anti-development is in




Fig. 8. Comparison of estimation results (MSE) between linear interpolation, geodesic interpolation, and classic Kalman filter with $\delta t = 0.2s$. The state to estimate is constant and the variance of the observation is $\sigma_w^2 = 1$.




Fig. 9. Cumulated error of estimation for $0 \leq t \leq 50$ for the linear interpolation, the geodesic interpolation, and the classic Kalman filter for different values of δt . The state to estimate is constant and the variance of the observation is $\sigma_w^2 = 1$. Each point represents the average of 20 simulations configured with the same parameters. For a sampling time larger than 0.45 (arbitrarily chosen), the performances are too poor to consider the filters to converge properly any-more.

one-to-one correspondance with the observation, the Kalman filter is used here as a reference filter (as it is known that it is not possible, in the mean square sense, to outperform it).

The results presented in Figure 7 have been obtained with $\delta t = 10^{-2}s$. This time step is small enough so that similar performances for the different methods of interpolation (geodesic and linear) as they should converge towards the same solution. However, differences start to appear if δt has larger values.

In Figure 8, δt is increased to $0.2s$. The variation in the observation time step changes the approximation of the process z_t and differences between the two interpolation methods (linear and geodesic) are visible. The linear interpolation is not as accurate as the geodesic interpolation. The linear approximation adds another source of error to the estimation, as a drawback to its simpler computational form.

Finally, in order to observe the influence of the sampling step δt on the performance of each interpolation method, Figure 9 illustrates the evolution of the cumulated error for $0 \leq t \leq 50$ at a fixed sample step for each method. Just like previously, the cumulated errors should be compared with the cumulated error induced by a proper Kalman filter. It appears that the geodesic interpolation does not create another error term despite that it is an approximation of $\delta z_{n\delta t}$. The cumulated error is the same as the Kalman filter one. The linear interpolation, however, is adding a supplementary error term. As the sampling step is increasing, the approximation is worse and worse. In the end, for $\delta t > 0.45s$, the sampling step is too high for any filter to perform a proper estimation of the state. For such cases, solutions such as the extended Kalman filter presented in [7] should be privileged.

VII. CONCLUSION

In this paper, a solution to the problem of filtering from partial observations is presented. The observed process takes its values in the Stiefel manifold while the signal of interest evolves on the rotation group. Due to the multiplicative nature of the noise, standard methods cannot be applied directly. A solution relying on the construction of an intermediate process, namely the anti-development, is proposed. This solution uses a Monte-Carlo method to overcome the problem of missing information. It is shown that the proposed algorithm allows, in certain contexts, to recover the whole set of components of the signal of interest despite the lack of observations. Finally, in the special case where the entire process can be observed, an optimal filter, together with interpolation methods, is given. This filter can be interpreted as a Kalman filter for observations on the rotation group.

REFERENCES

- [1] A. M. Bloch and J. E. Marsden, "Stabilization of rigid body dynamics by the energy-casimir method," *Systems and Control Letters*, vol. 14, no. 4, pp. 341 – 346, 1990.
- [2] A. Astolfi and M. Lovera, "Global spacecraft attitude control using magnetic actuators," in *American Control Conference, 2002. Proceedings of the 2002*, vol. 2, 2002, pp. 1331–1335 vol.2.
- [3] M. Zamani, J. Trumpf, and R. Mahony, "Near-optimal deterministic filtering on the rotation group," *Automatic Control, IEEE Transactions on*, vol. 56, no. 6, pp. 1411–1414, June 2011.
- [4] A. Barrau and S. Bonnabel, "Intrinsic filtering on $so(3)$ with discrete-time observations," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, Dec 2013, pp. 3255–3260.
- [5] M. Zamani, J. Trumpf, and R. Mahony, "Minimum-energy filtering for attitude estimation," *Automatic Control, IEEE Transactions on*, vol. 58, no. 11, pp. 2917–2921, Nov 2013.
- [6]
- [7] G. Bourmaud, R. Mégret, A. Giremus, and Y. Berthoumieu, "Discrete extended kalman filter on lie groups," *European Signal Processing Conference*, 2013.
- [8] P. Canet, "Kalman filter estimation of angular velocity and acceleration on-line implementation," Report, 1994.
- [9] A. Edelman, T. A. Arias, and S. T. Smith, "The geometry of algorithms with orthogonality constraints," *Siam J. Matrix Anal. Appl.*, vol. 20, no. 2, pp. 303–353, 1998.
- [10] O. Micka and A. J. Weiss, "Estimating frequencies of exponentials in noise using joint diagonalization," *Signal Processing, IEEE Transactions on*, vol. 47, no. 2, pp. 341–348, 1999.
- [11] Y. Chikuse, *Statistics on special manifolds*. Springer Lecture notes in Statistics, 2003.
- [12] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimisation algorithms on matrix manifolds*. Princeton University Press, 2008.
- [13] M. K. Ozdemir and H. Arslan, "Channel estimation for wireless ofdm systems," *Communications Surveys Tutorials, IEEE*, vol. 9, no. 2, pp. 18–48, 2007.
- [14] Y. M. Lui, J. R. Beveridge, and M. Kirby, "Canonical stiefel quotient and its application to generic face recognition in illumination spaces," in *Proceedings of the 3rd IEEE international conference on Biometrics: Theory, applications and systems*, ser. BTAS'09. IEEE Press, 2009, pp. 431–438.
- [15] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton, NJ: Princeton University Press, 2008.
- [16] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. Academic press, 1970.
- [17] A. Doucet and A. M. Johansen, "A tutorial on particle filtering and smoothing: fifteen years later," 2011.
- [18] M. Basseville and I. Nikiforov, *Detection of abrupt changes: Theory and application*, ser. Prentice hall information and system sciences series. Prentice Hall, 1993.